Correlation-Aware Distributed Caching and Coded Delivery

P. Hassanzadeh, A. Tulino, J. Llorca, E. Erkip

arXiv:1609.05836v1 [cs.IT] 19 Sep 2016

***Abstract*—Cache-aided coded multicast leverages side information at wireless edge caches to efficiently serve multiple groupcast demands via common multicast transmissions, leading to load reductions that are proportional to the aggregate cache size. However, the increasingly unpredictable and personalized nature of the content that users consume challenges the efficiency of existing caching-based solutions in which only *exact* content reuse is explored. This paper generalizes the cache-aided coded multicast problem to a source compression with distributed side information problem that specifically accounts for the correlation among the content files. It is shown how joint file compression during the caching and delivery phases can provide load reductions that go beyond those achieved with existing schemes. This is accomplished through a lower bound on the fundamental rate-memory trade-off as well as a correlation-aware achievable scheme, shown to significantly outperform state-of-the-art correlation-unaware solutions, while approaching the limiting rate-memory trade-off.**

I. Introduction and Setup

We consider a broadcast caching network with one sender (e.g., base station) connected to n receivers (e.g., access points or user devices) $\mathcal{U} = \{1, \dots, n\}$ via a shared error-free multicast link. The sender has access to a file library $\mathcal{F} = \{1, \dots, m\}$ composed of m files. Each file $f \in \mathcal{F}$ is represented by a vector of independent and identically distributed (i.i.d.) binary symbols of length F, $\mathsf{W}_f \in \mathbb{F}_2^F$, and therefore $H(\mathsf{W}_f) = F$ [1]. Each receiver has a cache of size MF bits for a real number $M \in [0, m]$, and receivers request files in an i.i.d. manner according to a demand distribution $\mathbf{q} = (q_1, \dots, q_m)$, where q_f denotes the probability of requesting file $f \in \mathcal{F}$. The file library realization is denoted by $\{W_f : f \in \mathcal{F}\}$, and we further assume that files can be correlated, i.e., $H(\mathsf{W}_{f_i}|\mathsf{W}_{f_j}) \leq F$, $\forall f_i, f_j \in \mathcal{F}$, and hence $H(\mathsf{W}_1, \dots, \mathsf{W}_m) \leq mF$. Such correlations are especially relevant among content files of the same category, such as episodes of a TV show or same-sport recordings, which, even if personalized, may share common backgrounds and scene objects. Hence, we assume that the joint distribution of the library files, denoted by $P_{\mathcal{W}}$, is not necessarily the product of the file marginal distributions.

P. Hassanzadeh and E. Erkip are with the ECE Department of New York University, Brooklyn, NY. Email: {ph990, elza}@nyu.edu

J. Llorca and A. Tulino are with Bell Labs, Nokia, Holmdel, NJ, USA. Email: {jaime.llorca, a.tulino}@nokia.com

A. Tulino is with the DIETI, University of Naples Federico II, Italy. Email: {antoniamaria.tulino}@unina.it

[1]The results derived in this paper can be easily extended to the case in which the symbols within each file are not necessarily binary and independent.

Consistent with the existing information-theoretic literature on cache-aided networks [1]–[5], we assume that the network operates in two phases: a caching (or placement) phase taking place at network setup, in which caches are populated with content from the library, followed by a delivery phase where the network is used repeatedly in order to satisfy receiver demands. The design of the caching and delivery phases forms what is referred to as a *caching scheme*.

The same network setup, but under the assumption of independent library files, was studied in [1]–[5]. These works characterized the order-optimal rate-memory trade-off for both worst-case [1], [2] and random demand settings [3]–[5]. In this context, each file in the library is treated as an independent piece of information, compressed up to its entropy F. During the caching phase, parts of the library files are cached at the receivers according to a properly designed *caching distribution*. The delivery phase consists of computing an *index code*, in which the sender compresses the set of requested files into a multicast codeword, only exploring perfect matches ("correlation one") among (parts of) the requested and the cached files, ignoring all other correlations that exist among the different (parts of) files.

The goal of this paper is to investigate the additional gains that can be obtained by exploring the correlations among the library content, in both caching and delivery phases. An initial step towards the characterization of the rate-memory region for a correlated library was given in [6]. The authors focus on lossy reconstruction in a scenario with only two receivers and one cache. In contrast, we address the more general multiple cache scenario, and focus on correlation-aware lossless reconstruction. We formulate the problem via information-theoretic tools where the caching phase allows for placement of arbitrary functions of a correlated library at the receivers, while the delivery phase becomes equivalent to a source coding problem with distributed side information. A correlation-aware scheme then consists of receivers storing content pieces based on their popularity as well as on their correlation with the rest of the library in the caching phase, and receiving compressed versions of the requested files according to the information distributed across the network and their joint statistics during the delivery phase. We introduce a lower bound for the rate-memory trade-off as well as an achievable scheme that includes a library compression step proceeding correlation-aware versions of the caching and delivery phases previously proposed in [3], [4]. We provide comparisons with state of art correlation-unaware schemes and an alternative correlation-aware achievable strategy introduced in [7].

The paper is organized as follows. Sec. II presents the information-theoretic problem formulation. Sec. III provides a lower bound on the fundamental rate-memory trade-off. An achievable scheme is described in Sec. IV, its performance is numerically validated in Sec. V, and the conclusion follows in Sec. VI.

II. Problem Formulation

Notation: For ease of exposition, we use $\{A_i\}$ to denote the set of elements $\{A_i : i \in \mathcal{I}\}$, with $\mathcal{I}$ being the domain of index i. $\mathbb{F}_2^*$ denotes the set of finite length binary sequences. We use $[x]^+$ to denote $\max\{x, 0\}$.

We consider the following information-theoretic formulation of a caching scheme for n receivers and m files:

- Initially, a realization of the library $\{W_f\}$ is revealed to the sender.
- **Cache Encoder:** At the sender, the cache encoder computes the content to be placed at the receiver caches by using the set of functions $\{Z_u : \mathbb{F}_2^{mF} \to \mathbb{F}_2^{MF} : u \in \mathcal{U}\}$, such that $Z_u(\{W_f\})$ is the content cached at receiver u, and M is the (normalized) cache capacity or alternatively memory. The cache configuration $\{Z_u\}$ is designed jointly across receivers, taking into account global system knowledge such as the number of receivers and their cache sizes, the number of files, their aggregate popularity, and their joint distribution $P_{\mathcal{W}}$. Computing $\{Z_u\}$ and populating the receiver caches constitutes the caching phase, which is assumed to happen during off-peak hours without consuming actual delivery rate.
- **Multicast Encoder:** Once the caches are populated, the network is repeatedly used for different demand realizations. At each use of the network, a random demand vector $\mathsf{f} = (\mathsf{f}_1, \ldots, \mathsf{f}_n) \in \mathcal{F}^n$ is revealed to the sender. We assume that f has i.i.d components distributed according to $\mathbf{q}$ and we denote by $\mathbf{f} = (f_1, \ldots, f_n)$ its realization. The multicast encoder is defined by a fixed-to-variable encoding function $X : \mathcal{F}^n \times \mathbb{F}_2^{mF} \times \mathbb{F}_2^{nMF} \to \mathbb{F}_2^*$, such that $X(\mathbf{f}, \{W_f\}, \{Z_u\})$ is the transmitted codeword generated according to demand realization $\mathbf{f}$, library realization $\{W_f\}$, cache configuration $\{Z_u\}$, and joint file distribution $P_{\mathcal{W}}$.
- **Multicast Decoders:** Each receiver $u \in \mathcal{U}$ recovers its requested file W_{f_u} using the received multicast codeword and its cached codeword, as $\widehat{W}_{f_u} = \rho_u(\mathbf{f}, X, Z_u)$, where $\rho_u : \mathcal{F}^n \times \mathbb{F}_2^* \times \mathbb{F}_2^{MF} \to \mathbb{F}_2^F$ denotes the decoding function of receiver u.

The worst-case (over the file library) probability of error of the corresponding caching scheme is defined as

$$P_e^{(F)} = \sup_{\{W_f : f \in \mathcal{F}\}} \mathbb{P}\left(\widehat{W}_{f_u} \neq W_{f_u}\right).$$

In line with previous work [3]–[5], the (average) rate of the overall caching scheme is defined as

$$R^{(F)} = \sup_{\{W_f : f \in \mathcal{F}\}} \frac{\mathbb{E}[J(X)]}{F}. \tag{1}$$

where $J(X)$ denotes the length (in bits) of the multicast codeword X.

Definition 1. *A rate-memory pair (R, M) is* achievable *if there exists a sequence of caching schemes for cache capacity (memory) M and increasing file size F such that $\lim_{F\to\infty} P_e^{(F)} = 0$, and $\limsup_{F\to\infty} R^{(F)} \leq R$.*

Definition 2. *The rate-memory region is the closure of the set of achievable rate-memory pairs (R, M). The rate-memory function $R(M)$ is the infimum of all rates R such that (R, M) is in the rate-memory region for memory M.*

In this paper, we find a lower bound and an upper bound on the rate-memory function $R(M)$, given in Theorems 1 and 2 respectively, and design a caching scheme (i.e., a cache encoder and a multicast encoder/decoder) that results in an achievable rate R close to the lower bound.

III. Lower bound

In this section, we derive a lower bound on the rate-memory function under uniform demand distribution using a cut-set bound argument on the broadcast caching-demand augmented graph [8]. To this end let $\mathcal{D}^{(j)}$ denote the set of demands with exactly j distinct requests.

Theorem 1. *For the broadcast caching network with n receivers, library size m, uniform demand distribution, and joint probability $P_{\mathcal{W}}$,*

$$R(M) \geq \liminf_{F\to\infty} \max_{\ell \in \{1,\cdots,\gamma\}} P_\ell \frac{[H(\{\mathsf{W}_f : f \in \mathcal{F}\}) - \ell MF]^+}{\left\lfloor \frac{m}{\ell} \right\rfloor F}$$

where $\gamma = \min\{n, m\}$, $P_\ell = P(\mathbf{d} \in \bigcup_{j\geq\ell} \mathcal{D}^{(j)})$ and $H(\{\mathsf{W}_f : f \in \mathcal{F}\})$ is the entropy of the entire library.

Proof: The proof of Theorem 1 is given in Appendix A. ∎

IV. Library Compressed Cache-Aided Coded Multicast (Comp-CACM)

In this section we describe a correlation-aware caching scheme, termed Comp-CACM. The proposed Comp-CACM scheme consists of a *Correlation-Aware Library Compressor*, a *Random Fractional Cache Encoder* and a *Coded Multicast Encoder*. In this scheme:

- *Step.1*: The library files are divided into two subsets, *I-files* and *P-files*.
- *Step.2*: The P-files are conditionally compressed with respect to the I-files.
- *Step.3*: The compressed files are divided into packets, and a portion of these packets is cached at each receiver according to a caching distribution.
- *Step.4*: During the delivery phase, the receivers requesting an I-file receive a linear combination of the uncached packets of the corresponding I-file, and those requesting a P-file receive a linear combination of the uncached packets of the corresponding P-file and the uncached packets of the associated I-file.

As described in more detail in the subsequent sections, Steps 1 and 2 are carried out by the Correlation-Aware Library

Compressor, Step 3 is done by the Random Fractional Cache Encoder and the Coded Multicast Encoder performs Step 4. Differently from existing correlation-unaware schemes, in Comp-CACM, the cache encoder operates on a compressed version of the library obtained by the Correlation-Aware Library Compressor. Even though the delivery phase is still based on computing an index code, this phase needs to be properly modified to account for the fact that the content stored in the receiver caches is a compressed version of the library files.

A. Correlation-Aware Library Compressor

In Steps 1 and 2, prior to passing the library to the cache encoder, the correlations among the library content are exploited to compress the files according to their joint distribution $P_{\mathcal{W}}$ as well as their popularity, resulting in a library composed of *inter-compressed* files.

Definition 3. *(δ-**Correlated Files**) For a given threshold $\delta \leq 1$ and file size F, we say file W_f is δ-correlated with file $\mathsf{W}_{f'}$ if $H(\mathsf{W}_f, \mathsf{W}_{f'}) \leq (1+\delta)F$ bits.*

Using the above notion of δ-correlated files, the file library $\mathcal{F} = \mathcal{F}_I \cup \mathcal{F}_P$ is partitioned into I-files $\mathcal{F}_I$, and inter-compressed P-files $\mathcal{F}_P$, as follows:

1) Initialize the file set to the original library $\mathcal{F}$.
2) Compute the *aggregate popularity* for each file in the file set as the probability that at least one receiver requests at least one file from its δ-*file-ensemble*, defined as the given file and all its δ-correlated files in the file set.
3) Choose the file with the highest aggregate popularity, W_{f^*}, and assign it to the I-file set $\mathcal{F}_I \subseteq \mathcal{F}$. All other files in the δ-file-ensemble of W_{f^*} are inter-compressed using W_{f^*} as reference and assigned to the P-file set $\mathcal{F}_P \subseteq \mathcal{F}$.
4) Update the file set by removing all files contained in the δ-file-ensemble of W_{f^*}.
5) Repeat steps 2-4 until the file set is empty.

The partitioning procedure described above results in a compressed library, $\{\overline{W}_f : f \in \mathcal{F}\}$, composed of I-files $\mathcal{F}_I$ each with entropy F, and P-files $\mathcal{F}_P$ inter-compressed with respect to their corresponding I-files, with conditional entropy less than δF. The value of δ will be further optimized as a function of the system parameters $(n, m, M, \mathbf{q}, \mathcal{P}_W)$ in order to minimize the (average) rate of the overall proposed caching scheme given in (1).

The following example illustrates the compression of the library for a simple system setup.

Example 1. *Consider a library consisting of $m = 4$ uniformly popular files $\{W_1, W_2, W_3, W_4\}$ each with entropy F bits. We assume that the file pairs $\{W_1, W_2\}$ and $\{W_3, W_4\}$ are independent, while correlations exist between W_1 and W_2, and between W_3 and W_4. Specifically, $H(\mathsf{W}_1|\mathsf{W}_2) = H(\mathsf{W}_2|\mathsf{W}_1) = F/4$ and $H(\mathsf{W}_3|\mathsf{W}_4) = H(\mathsf{W}_4|\mathsf{W}_3) = F/4$. Since files are equally popular and each file is correlated with one other file, each file has the same aggregate popularity. We choose file W_2 as the first I-file and assign its δ-correlated file, W_1, to the P-file set. The file set is then updated to $\{W_3, W_4\}$. Again, either file can be assigned to the I-file set as they have the same aggregate popularity. We choose W_4 as the second I-file and assign W_3 to the P-file set. This results in the compressed library $\{\overline{W}_f\} = \{\overline{W}_1, W_2, \overline{W}_3, W_4\}$ with $\mathcal{F}_I = \{2, 4\}$ and $\mathcal{F}_P = \{1, 3\}$. $\overline{W}_1(\overline{W}_3)$ denotes P-file $W_1(W_3)$ inter-compressed with respect to I-file $W_2(W_4)$.*

B. Random Fractional Cache Encoder

After compressing the library, in Step 3, Comp-CACM applies a slightly modified version of the random fractional cache encoder introduced in [3], [4] to the compressed library.

Specifically, given the compressed library $\{\overline{W}_f\}$, the cache encoder divides each file into equal-size packets, such that each file $\overline{W}_f$ is divided into $H(\overline{W}_f)/b \in \mathbb{N}$ packets of length b bits. The packet length b is chosen such that all files in the compressed library (of possibly different lengths) are divided into an integer number of packets. The sender then populates the caches by randomly selecting and storing $p_f M$ fraction of the packets of each file $\overline{W}_f$,$f \in \mathcal{F}$ at each receiver, where $\mathbf{p} = (p_1, \ldots, p_m)$, referred to as the *caching distribution*, is a vector with elements $0 \leq p_f \leq 1/M$, satisfying the cache constraint as $\sum_{f\in\mathcal{F}_I} p_f + \delta \sum_{f\in\mathcal{F}_P} p_f = 1$.[2] The caching distribution is optimally designed to minimize the rate of the corresponding correlation-aware delivery scheme as expressed in (1) while taking into account global system parameters $(n, m, M, \mathbf{q}, P_{\mathcal{W}})$. The rate expression in (1) is normalized to the file size, and is independent of packet length b when $F \to \infty$. In the following, we denote the i^{th} packet of file $W_f, f \in \mathcal{F}$ by W_{fi}.

Example 2. *In Example 1 further assume that the sender is connected to $n = 2$ receivers $\{u_1, u_2\}$ each with a cache size $M = 1$. For the compressed library $\{\overline{W}_1, W_2, \overline{W}_3, W_4\}$ described in Example 1, we assume packet length $b = F/4$ and a caching distribution such that $p_2 = p_4 = 1/4$ and $p_1 = p_3 = 1$. This corresponds to splitting files W_2 and W_4 into 4 packets, each with entropy $F/4$, and randomly storing 1 packet from each file at each receiver. Files $\overline{W}_1$ and $\overline{W}_3$ comprise one packet each and are both cached at the receivers. We assume that packets $\{\overline{W}_1, W_{21}, \overline{W}_3, W_{41}\}$ are cached at u_1 and $\{\overline{W}_1, W_{22}, \overline{W}_3, W_{42}\}$ are cached at u_2, as shown in Fig 1.*

C. Coded Multicast Encoder

The goal of the multicast encoder of Step 4 is to generate a coded multicast codeword that allows each receiver to losslessly reconstruct its requested file, given that the cached content and the resulting conflict graph is composed of inter-compressed file pieces. Each receiver recovers its demand using the multicast codeword and its cache content. The multicast codeword can be a linear combination of the

[2] As in [4], the caching distribution defines the number of packets cached from each file; however, since here I-files and inter-compressed P-files have different lengths and are therefore divided into a different number of packets, the cache constraint differs from that of [4].

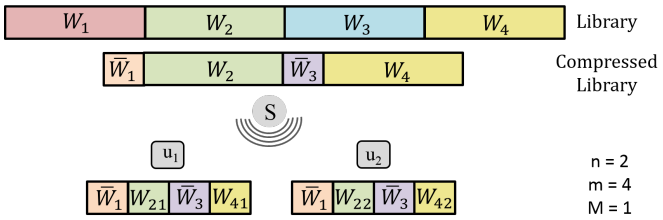

Fig. 1: Network setup and cache configuration of Examples 1-3.

packets needed by each receiver for demand recovery. The objective is to find the minimum multicast transmissions (termed as the multicast rate) given a set of demands and the cache content of all receivers. The problem can be formulated as the coloring of the corresponding conflict graph. However, differently from the classical setting described in [3]–[5], where each requested packet corresponds to a vertex in the conflict graph, here the vertices are composed of: 1) for each receiver requesting an I-file, the uncached packets of the corresponding I-file, and 2) for each receiver requesting a P-file, the uncached packets of the corresponding P-file and the uncached packets of the associated I-file [9]. As in the classical setting, a directed edge between two vertices of the graph means that one receiver's desired packet interferes with some other receiver's desired packet (i.e is not cached at that receiver).

Example 3. *For the cache configuration given in Example 2, and demand realization* $\mathbf{f} = (1, 2)$*, the sender multicasts* $W_{21} \oplus W_{22}, W_{23}, W_{24}$*, which corresponds to a total rate of* $R = 0.75$*. The transmitted codeword enables the reconstruction of file* W_2 *at* u_1*, and with* $\overline{W}_1$ *being in its cache, receiver* u_1 *is able to losslessly reconstruct* W_1*. Receiver* u_2 *is able to reconstruct* W_2 *using the received packets. However, a correlation-unaware scheme (e.g., [2], [3]) would first compress the files separately and then cache* $1/4^{th}$ *of each file at each receiver, which would result in a total rate* $R = 1.25$ *regardless of the demand realization [2], [3].*

D. Rate Upper Bound for Comp-CACM

In this section, we upper bound the rate-memory function given in Definition 2 by evaluating the performance of the Comp-CACM under the uniform demand scenario.[3] In stating the theorem, we assume that the library file correlation is such that for a given δ, there are at least κ files in the δ-file-ensemble of each file $f \in \mathcal{F}$. In other words, κ, which is a function of δ and of the joint library distribution $P_{\mathcal{W}}$, is the largest value such that for each file $f \in \mathcal{F}$, there are at least κ files δ-correlated with f.

Theorem 2. *Consider a broadcast caching network with* n *receivers, library size* m*, and uniform demand distribution. Then, the rate-memory function,* $R(M)$*, is upper bounded as*

$$R(M) \leq \inf_{\delta} \left\{ \min\{\psi(\delta, p_I^*, p_P^*, M), \bar{m}\} \right\}, \tag{2}$$

where

$$\begin{aligned}
&\psi(\delta, p_I, p_P, M) = \sum_{\ell=1}^{n} \binom{n}{\ell} \Big(\lambda_\ell(p_I) + \delta \lambda_\ell(p_P) \Big), \\
&\{p_I^*, p_P^*\} = \arg\min_{p_I, p_P} \{\psi(\delta, p_I, p_P, M), \bar{m}\} \\
&\qquad \text{s.t.} \quad p_I + \delta(\kappa - 1) p_P = \frac{\kappa}{m} \\
&\bar{m} = m\Big(1 - \Big(1 - \frac{1}{m}\Big)^n\Big),
\end{aligned} \tag{3}$$

with

$$\lambda_\ell(p) \triangleq (pM)^{(\ell-1)} (1 - pM)^{(n-\ell+1)}.$$

with κ *in (3) being a function of* δ *and* $P_{\mathcal{W}}$*.*

Proof: The proof of Theorem 2 is given in [9]. ■

Note that for a given δ, the argument of the infimum on the right-hand side of (2) represents an upper bound on the Comp-CACM rate achieved with the caching distribution optimized for uniformly popular files, given by $p_f = p_I^*$, $\forall f \in \mathcal{F}_I$, and $p_f = p_P^*$, $\forall f \in \mathcal{F}_P$. As mentioned in the previous section, the caching distribution of Comp-CACM is designed to minimize the rate of the corresponding correlation-aware delivery scheme. Consequently, the caching distribution of Comp-CACM not only depends on the demand distribution, as in correlation-unaware schemes, but also depends on the correlations among the library content. Hence, under a uniform demand distribution, differently from correlation-unaware schemes, where all files are uniformly cached due to their uniform popularity, the Comp-CACM caching distribution assigns a larger portion of the memory to storing I-files compared to storing P-files, due to I-files having a higher *aggregate popularity* than P-files (see Sec. IV-A).

V. Numerical Results and Discussions

We numerically compare the rates achieved by the proposed correlation-aware scheme Comp-CACM, with respect to the following schemes: Local Caching with Unicast Delivery (LC/U), Local Caching with Naive Multicast Delivery (LC-NM), Random Popularity based (RAP) Caching with Coded Multicast Delivery (RAP/CM) [4], all of which are state-of-the-art caching schemes that do not consider and exploit content correlation. We also compare the rates with an alternative correlation-aware scheme, termed CA-RAP/CM introduced in [7], [9], in which receivers store content pieces from the original (not inter-compressed) library based on their popularity and their correlation with the rest of the file library during the caching phase, and receive compressed versions of the requested files according to the information distributed across the network and their joint statistics during the delivery phase. Even though the caching of content in CA-RAP/CM is not done as efficiently as in Comp-CACM, the delivery is more efficient. All of the schemes mentioned

[3]Closed-form expressions for the achievable rate of Comp-CACM for the general setting described in Sec. I are given in [9].

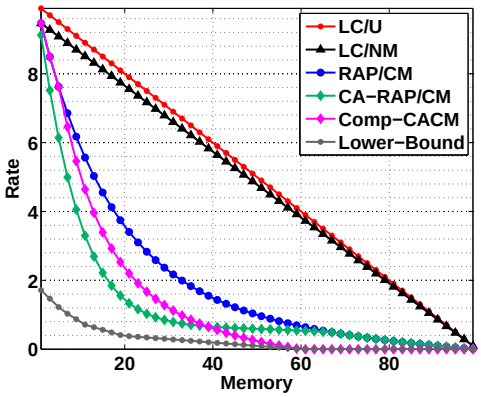

Fig. 2: Comparison of Comp-CACM and the lower bound in Theorem 1 with other caching schemes in a network with uniform demand distribution, $n = 10$, $m = 100$, $\delta = 0.2$, and $\kappa = 2$.

above are compared with the rate-memory function lower bound given in Theorem 1.

We consider a broadcast caching network with $n = 10$ receivers requesting files according to a uniform demand distribution from a library composed of $m = 100$ files, in which there are κ files in the δ-ensemble of each file.

Fig. 2 displays the rate-memory trade-off as the memory size varies from 0 to 100 files for $\delta = 0.2$ and $\kappa = 2$. The figure plots the expected achievable rate R, in number of transmissions normalized to the file size, versus memory size (cache capacity) M, normalized to the file size. As expected, the correlation-aware schemes outperform schemes that are oblivious to the file correlations. Comp-CACM achieves a $3.5\times$ reduction in the expected rate compared to LC-U and a $1.5\times$ reduction compared to RAP/CM for $M = 20$. When comparing Comp-CACM and CA-RAP/CM, for small memory size M, CA-RAP/CM outperforms Comp-CACM, while the opposite is observed for large M. This indicates that joint compression of the file library has an especially relevant effect at large memory size, since each receiver is able to cache most of the compressed library. In fact, Comp-CACM achieves zero rate at around $M = 60$, indicating that every receiver is able to cache the entire library in its compressed form. For small cache size M, however, the extra bits that Comp-CACM is forced to request in the delivery phase due to the initial library compression undermines its benefit with respect to CA-RAP/CM.

VI. Conclusion

In this paper we have formulated the caching problem using information-theoretic tools and have provided a lower bound on the fundamental rate-memory trade-off. We have proposed Comp-CACM, an achievable correlation-aware scheme, in which, prior to storing the content files, the library is jointly compressed while taking into account the popularity of the files as well as their joint distribution. This leads to caching of the more relevant content during the caching phase, which, with the transmissions during the delivery phase that ensure lossless reconstruction of the demand, results in the more efficient delivery of content. We have then compared the performance of Comp-CACM with CA/RAP-CM, another correlation-aware scheme, as well as with the conventional caching schemes.

Part of our ongoing work focuses on i) designing improved achievable schemes that incorporate the beneficial compression aspects of Comp-CACM into CA-RAP/CM, such as jointly compressing the cache configuration at each receiver, and ii) studying the effect of relevant system parameters that characterize the correlated library such as δ and κ.

Appendix A
Proof of Theorem 1

Denoting by $\mathcal{D}^{(j)}$ the set of all demands that contain at least j distinct requests, it is immediate to prove that

$$R(M) \;\geq\; P\left(\mathbf{d} \in \bigcup\nolimits_{j\geq \ell} \mathcal{D}^{(j)}\right) R_\ell(M) \qquad (4)$$

where $R_\ell(M)$ is the average rate needed to satisfy the user demands containing exactly ℓ distinct requests. Due to the uniform popularity assumption, we have that

$$R_\ell(M) = \frac{1}{|\mathcal{D}^{(\ell)}|} \sum_{\mathbf{d}\in\mathcal{D}^{(\ell)}} R_\ell(M,\mathbf{d}) \qquad (5)$$

where $R_\ell(M,\mathbf{d})$ denotes the rate needed to satisfy user demand $\mathbf{d}$. Let $\{\mathbf{d}_1,\ldots,\mathbf{d}_{\lfloor \frac{m}{\ell} \rfloor}\}$ denote a set of $\lfloor \frac{m}{\ell} \rfloor$ demands in $\mathcal{D}^{(\ell)}$ such that the union of all their requested files is equal to the entire library. Furthermore, by the information-theoretic cut-set bound we have that

$$\sum_{\mathbf{d}\in\left\{\mathbf{d}_1,\ldots,\mathbf{d}_{\lfloor \frac{m}{\ell} \rfloor}\right\}} R_\ell(M,\mathbf{d})F + \ell M F \geq H(\{\mathsf{W}_f : f \in \mathcal{F}\}) \qquad (6)$$

Replacing (6) in (4), Theorem 1 follows.

References

[1] M. A. Maddah-Ali and U. Niesen, "Fundamental limits of caching," *IEEE Transactions on Information Theory*, vol. 60, no. 5, pp. 2856–2867, 2014.

[2] M. Maddah-Ali and U. Niesen, "Decentralized coded caching attains order-optimal memory-rate tradeoff," *IEEE/ACM Transactions on Networking*, no. 99, pp. 1–8, 2014.

[3] M. Ji, A. Tulino, J. Llorca, and G. Caire, "On the average performance of caching and coded multicasting with random demands," in *Proc. IEEE International Symposium on Wireless Communications Systems (ISWCS)*, 2014, pp. 922–926.

[4] ——, "Order-optimal rate of caching and coded multicasting with random demands," *arXiv:1502.03124*, 2015.

[5] M. Ji, A. M. Tulino, J. Llorca, and G. Caire, "Caching and coded multicasting: Multiple groupcast index coding," in *Proc. IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2014, pp. 881–885.

[6] R. Timo, S. S. Bidokthi, M. Wigger, and B. Geiger, "A rate-distortion approach to caching," *preprint http://roytimo. wordpress. com/pub*, 2015.

[7] P. Hassanzadeh, A. Tulino, J. Llorca, and E. Erkip, "Cache-aided coded multicast for correlated sources," *in Proc. IEEE International Symposium on Turbo Codes and Iterative Information Processing (ISTC)*, 2016.

[8] J. Llorca, A. M. Tulino, K. Guan, and D. Kilper, "Network-coded caching-aided multicast for efficient content delivery," in *Proc. IEEE International Conference on Communications (ICC)*, 2013, pp. 3557–3562.

[9] P. Hassanzadeh, A. Tulino, J. Llorca, and E. Erkip, "Rate-memory trade-off for caching and delivery of correlated sources," *preprint http://engineering.nyu.edu/eelab/people/Parisa*.